

Certificate of Mailing/Facsimile 37 CFR §1.8(a)

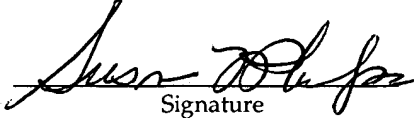
I hereby certify that this correspondence is being:

X deposited with the United States Postal
Service as first class mail in an envelope with
sufficient postage addressed to the:

_____ facsimile transmitted to [Fax Number] to the:

Mail Stop Appeal Brief - Patents
Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on June 14, 2005.

Susan L. Phelps


Signature

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Applicant: Blackmore et al.	:	GROUP ART UNIT: 2153
Serial No.: 09/619,051	:	Examiner: A. Choudhary
Filed: July 18, 2000	:	June 14, 2005
Title: Mechanisms for Efficient Message:	:	Lawrence D. Cutter
Passing with Copy Avoidance in a	:	IBM Corporation
Distributed System Using Advanced	:	2455 South Road, M/S P386
Network Devices	:	Poughkeepsie, NY 12601

SUBSTITUTE APPEAL BRIEF UNDER 37 C. F. R. § 1.192

Mail Stop Appeal Brief-Patents
Hon. Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an appeal from a final rejection dated September 17, 2004, rejecting all of the claims in the above-identified patent application. This is a substitute Appeal Brief submitted in response to the Examiner's assertion of defects in formality of the original Brief submitted on February 17, 2005. The prior Brief was accompanied by a transmittal letter authorizing the charging of appellants' deposit account for payment of the requisite fee set forth in 37 CFR § 1.17(c).

I. REAL PARTY IN INTEREST

This application is assigned to International Business Machines Corporation by virtue of an assignment executed on July 18, 2000, and also on November 2, 2000 by the co-inventors and recorded on November 6, 2000, in the U.S. Patent and Trademark Office on Reel 011239 and Frame 0948. Therefore, the real party in interest is International Business Machines Corporation.

II. RELATED APPEALS AND INTERFERENCES

To the knowledge of the Appellants, Appellants' undersigned legal representative and the Assignee, there are no other appeals or interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the instant appeal.

III. STATUS OF CLAIMS

Claims present: 1-4;
Claims allowed: none;
Claims rejected: 1-4;
Claims objected to: none;
Claims cancelled: 1 herein.
Claims appealed: 1-3.

At present, claims 1-4 stand rejected under 35 U.S.C. § 103 based upon the U.S. Patent to Sethuram et al. (U.S. Patent No. 5,828, 903 issued on October 27, 1998, and having a priority date based upon a continuing application filed on September 30, 1994). Applicants' claims 1-4 are rejected based on this patent and 35 U.S.C. § 102. Claim 4 is cancelled herein.

In response to a final rejection dated September 17, 2004, Applicants submitted a response under 37 CFR § 1.116 on November 16, 2004. An Advisory Action was sent on

March 1, 2005, but at a time after applicants' Appeal Brief had been submitted. In a discussion with the Examiner on December 16, 2004, the Examiner indicated that she had not even seen the submitted response under 37 CFR § 1.116. Accordingly, on December 16, 2004, a Notice of Appeal was filed. This brief is hereby being submitted as a result of the filing of the aforementioned Notice of Appeal and in response to the Notification of Non-Compliant Appeal Brief with certain formalities which have been addressed hereby.

Accordingly, Applicants are appealing the rejection of claims 1-3.

IV. STATUS OF AMENDMENTS

All amendments in the present application have been entered and are reflected in the Claims Appendix in section VIII herein.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present invention is directed to an efficient mechanism for message passing in a data processing system by avoiding the use of unnecessary message copies. More particularly, the present invention is directed to the use of an interface mechanism which efficiently implements zero-copy transport protocols. The overhead associated with unnecessary message copying is eliminated in the present invention through the use of copy avoiding message passing protocols. The interfaces provided in the present invention allow a user to prepare the source and target buffers and to enable the network device adapter to directly transfer, via direct memory access operations, data from an adapter directly into desired target locations. More particularly, with respect to the present invention on the target/receiving side, the receiver allows the network device to transfer, via the DMA mechanism, incoming data directly into target memory locations rather than into an intermediate buffer. It is important to note that the present invention provides a dynamic association to desired target memory locations. In contrast, it is seen that the cited art teaches only a dynamic association with an intermediate buffer.

As indicated in applicants' specification, the phrase "desired target memory location" refers to locations specified by a user, that is, by application level programming. In this regard, attention is directed to page 6, lines 15-18, wherein it is stated that it is an object of the claimed invention "to enable users running applications in their own address spaces on one data processing system to be able to transfer data . . . into a user's address space in another data processing system" Thus, when one naturally asks, "Desired by what or whom?", the answer is: that which is desired by an application level user. Hence, applicants treat the phrases "desired target memory" and "user target memory" as being synonymous.

For example, applicants describe the protocol for a zero copy transport as including the following steps:

- (1) The user prepares the send and receive buffers (page 9, line 8);
- (2) A mutually agreed upon tag is associated with the buffer (page 9, line 13);
- (3) The prepared buffer is posted to the communication layer (page 9, line 17);
and
- (4) The user issues a send operation that specifies . . . the mutually agreed upon tag that provides a mapping to the receive buffer at the target . . . (page 9, line 21).

These aspects of the invention are pointed out in support of applicants' position that it is the user that establishes the desired data destination, that is, the data target memory location desired by the user.

More particularly, in applicants' claim 1, a method is provided for sending a message from the memory (103) of a first data processing system (101) to the memory (203) of a second data processing system (201). See also Figure 7, items 1-5. The process employs the transmission of the message from the first system (201) to a temporary memory (152) in a

communications adapter (205) coupled to the second system (201). See page 21, lines 8-10. An indication that the temporary memory (152) contains the message is sent to the second system (201) from its associated communications adapter (205). See page 212, lines 10-11. Real address information indicating desired target memory locations is transferred from the second system (201) to its adapter (205). See Figure 5 and page 21, lines 2-4. The message from temporary adapter memory 152 is transferred directly into target memory locations in the memory 203 of the second system 201 via a DMA (Direct Memory Access) operation. See page 19, lines 13-16. An indication that the target locations now contain the message received is sent to the second system (201). See item 2 in Figure 7 and page 22, lines 16-18. Lastly, an acknowledgment of receipt of the message is sent from the second system (201) to the first system (101). See reference numeral 177 in Figure 7 and page 22, lines 18-20.

Claim 3 is similar to claim 1 except that it includes an additional step of establishing an association between the message and real address information indicating desired target memory locations. See page 9, lines 13-16, and page 21, lines 11-15.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Whether claims 1-3 are anticipated by the patent to Sethuram et al. (that is, rejectable under 35 USC § 102.)

VII. ARGUMENTS

Applicants submit that there are two arguments that are applicable in support of their request for withdrawal of the rejection.

A. User Directed Target Determination

In order to fully appreciate the significant differences between Applicants' claimed invention and the teachings found in the cited patent to Sethuram et al., the Examiner's attention is directed to column 2, lines 22-25:

"Virtual registers (VRs) are established in the adapter local memory to identify the location of buffers in the host memory used to store data transferred to and from the network." (Emphasis added herein.)

From the quotation above, it is clear that the DMA operation performed by Sethuram et al. is one in which transmitted data is stored within buffers in the receiving processor. In contrast, Applicants' claimed process specifically recites a transfer of the message from the adapter directly into target memory locations. These locations are the locations in the receiving target processor as employed by applications running on the target processor. Applicants' process is enabled to work in this fashion as a result of the fact that the third step in Applicants' recited process transfers real address information indicating desired target memory locations from the receiving processor to the adapter. This real address information, provided from the target processor, is supplied to the adapter. This transfer of real address information allows data to be written directly to target memory locations.

In stark contrast, it is seen that the teachings of Sethuram et al. require the writing of received message data into buffer locations in the target processor. Further operations are required in the process of Sethuram et al. to transfer the received data into actual target locations. Sethuram et al. operate only to provide transfer into buffers identified to the network not to application programs running on the relevant processors. Accordingly, it is seen that the process recited in Sethuram et al. is not in any sense the zero-copy process that is taught by the present applicants. When Sethuram et al. refer to transferring data directly to into the host memory, it is clear that they are referring to a transfer into an intermediary buffer in the host memory. Sethuram et al. do not teach, disclose or suggest a mechanism by which DMA transfers are made directly into specific target memory locations. In particular, Sethuram et al. do not teach, disclose or suggest that such target memory locations are provided via the transfer of real address information from the receiving target host to the adapter. While Sethuram et al. appear to indicate that virtual register information containing real address data may be established with association to certain virtual circuits, either dynamically or ahead of time, nowhere is there any teaching, disclosure or suggestion that these virtual registers contain real address information that is supplied directly from the target processor or that such real address information points to actual desired target memory locations. Rather it is seen that the teachings of Sethuram et al. are

that such transfers are only to intermediate memory buffers in receiving host, not to identified ultimate target locations. These buffers are known to the operating system but they are not directly accessible by end users or application level programs. As long as transfers are made to intermediate buffers, the transfer is not a zero-copy transfer and suffers the inefficiencies taught against by the present applicants.

The non zero-copy aspects from the patent to Sethuram et al. are also apparent on their transmittal side. In this regard, the Examiner's attention is directed to column 2, lines 33-36 which state as follows:

"When data is to be transmitted, the host device writes the data into the host memory in the buffer identified for a specified virtual circuit and the adapter is notified of the data to be transmitted."

From the above, it is seen that when messages are transmitted, the message is written (from an application's address space) first into a host memory buffer that is associated with a specified virtual circuit. Accordingly, it is seen that, as on the receiving side, the transmitting side associates a buffer with a virtual circuit. Even if the virtual registers contain real address information (currently a point of difference between Applicants and the Examiner), the real addresses point to intermediate buffer storage, not to ultimate target locations. The patent to Sethuram et al. is therefore seen to be a buffer-to-buffer transfer, not a user-to-user transfer.

B. Transmission of an Acknowledgment

It is furthermore noted that the rejection of Applicants' claims 1-3 is based upon 35 U.S.C. § 102. This is a narrow ground of rejection requiring each and every recited claim element to be found within the four corners of a single cited document. Here, the Examiner relies upon the patent to Sethuram et al. as that underlying document. However, it is specifically noted that Applicants' claims 1-3 include the recitation of a step in which an acknowledgment is transmitted from the second or receiving data processing system back to the first or the transmitting data processing system (the sixth and last recited claim step in Applicants' claim 1, for example). While Sethuram et al. appear to provide some notification between the receiving

processing system and its associated adapter (column 6, lines 51-56), there is no teaching, disclosure or suggestion in Sethuram et al. for the transmission of the acknowledgment which is specifically recited in Applicants' claims 1-3, that is, transmission of an acknowledgment back to the sending processor. It is seen that on this basis alone, the rejection under 35 U.S.C. § 102 cannot be sustained.

In pointing out the differences between Applicants' claimed invention and the teachings of Sethuram et al., it is to be specifically noted that Applicants' claims refer to a "desired target memory location." In contrast, the patent to Sethuram et al. refers only to buffers. A buffer is not the same thing as a desired target location. In Applicants' claimed invention, the data is written to the ultimate destination, not to an intermediate buffer even if that buffer is in the host memory. In Sethuram et al., applications using the transmitted data are still required to extract this information from the buffer. In contrast, in Applicants' claimed invention, desired target memory locations are ultimate destinations. In this regard, the Examiner's attention is directed to page 4, lines 5-7 of Applicants' specification wherein it is stated:

"Communication software is invoked so that the CPU is not engaged in staging the message through intermediate buffers in the communication software. The interfaces of the present invention allow the user to prepare (pin and map) the source/target buffer and enable the network device (adapter) to DMA (direct memory access) directly from the user buffer into the network. On the target/receiving side, the receiver posts user buffers to the network device so as to allow the network device to transfer, via a DMA mechanism, incoming data directly into the user target buffer."

Accordingly, it is seen that in Applicants' claimed invention, transfer is directly into memory locations controlled by users, not into a communications buffer controlled by the operating system. This is indeed true zero-copy transfer.

Therefore, it is seen that Applicants' claim steps 3 and 4 in claim 1 indicate that message data is transferred directly into desired target memory locations. In contrast, it is seen that the patent to Sethuram et al. teaches only the transfer of such data into intermediate communication buffers.

Accordingly, it is seen that there are indeed significant differences between that which is recited in Applicants' claims and that which is taught by Sethuram et al. These differences are in addition to the differences pointed out in Applicants' previously submitted response under 37 CFR § 1.116. Accordingly, for all of the indicated reasons, it is seen that since there are significant differences between that which is claimed and that which is taught by the cited art, the rejection of Applicants' claims 1-3 under 35 USC § 102 based upon the patent to Sethuram et al. cannot be supported. It is therefore respectfully requested that the rejection of Applicants' claims 1-3 be reversed.

VIII. CLAIMS APPENDIX

1. A method for sending a message stored in the memory of a first data processing system into the memory of a second data processing system, said method comprising the steps of:

transmitting said message from said first data processing system to a temporary memory in an adapter which is connected to said second data processing system;

transferring, from said adapter to said second data processing system, an indication that the temporary memory in said adapter contains the message received from said first data processing system;

transferring, from said second data processing system to said adapter, real address information indicating desired target memory location for said message;

transferring said message, from said temporary memory in said adapter, directly into target memory locations in the memory of said second data processing system, said transfer occurring via direct memory access;

transferring, from said adapter to said second data processing system, an indication that said target locations now contain the message received from said first data processing system; and

transmitting an acknowledgment of receipt of said message from said second data processing system to said first data processing system.

2. The method of Claim 1 further including the step of advancing indicators in said first data processing system in preparation of transmitting another message, whereby a number of messages may be sent in rapid sequence.

3. A method for sending a message from a first data processing system to the memory of a second data processing system, said method comprising the steps of:

establishing an association between said message and real address information indicating desired target memory locations for said memory;

transmitting said message from said first data processing system to a temporary memory in an adapter which is connected to said second data processing system;

transferring, from said adapter to said second data processing system, an indication that said temporary memory in said adapter contains the message received from said first data processing system;

transferring, from said second data processing system to said adapter, said real address information;

transferring said message, from said temporary memory in said adapter, directly into said target memory locations in the memory of said second data processing system, said transfer occurring via direct memory access;

transferring, from said adapter to said second data processing system an indication that said target locations contain the message received from said first data processing systems; and

transmitting an acknowledgment of receipt of said message from said second data processing system to said first data processing system.

4. (Cancelled)

IX. EVIDENCE APPENDIX

Not applicable.

X. RELATED PROCEEDINGS APPENDIX

None.

RESPECTFULLY SUBMITTED

Date: June 14, 2005

By: Lawrence D. Cutter

Lawrence D. Cutter, Senior Attorney

Registration No. 28,501

Phone: (845) 433-1172

FAX: (845) 432-9786